Staying Ahead of the Power Curve

Q&A with Intel Senior Fellow Mark T. Bohr

Table of Contents

(Click on page number to jump to sections)

Staying Ahead of the Power Curve

Silicon manufacturers have been generating a lot of news lately about energy-efficient platforms. Energy-efficient platforms begin with energy-efficient transistors. In order to learn more about Intel's advanced transistors we spoke with Mark T. Bohr, Intel Senior Fellow and Director of Process Architecture and Integration for Intel's Technology and Manufacturing Group, about how Intel silicon innovations are helping deliver platforms with unprecedented levels of energy-efficient performance.

Energy-Efficient Performance Starts at the Silicon Level

We've seen press lately about how Intel is "back in the lead" on energy-efficient performance with its processor technology. How do you see Intel's position in processor technology, and how are the latest developments in silicon process contributing to Intel's energy-efficient platforms strategy?

First of all, it's important to realize that energy-efficient performance is not something new. It is fundamental to what Intel has been doing for years. One of the most recent and hugely successful examples is Intel® Centrino® mobile technology, a comprehensive platform approach that brought new levels of performance and energy efficiency to mobile computing. To achieve even greater levels of energy-efficient performance, we have taken a comprehensive approach to building energy-efficient processors and platforms for all market sectors, and that approach starts at the silicon level.

As far as "taking back" our leadership, I think we've always been the leaders at the silicon level. For example, we are at least a year ahead of the competition with our 65-nanometer (nm) process. While competitors are talking about "sampling" or "product or process qualification" or "initial production," Intel has already shipped tens of millions of 65nm products. The Intel® Core™ Duo processor, Intel® Core™2 Duo processor, and the Intel® Pentium® processor Extreme Edition 955 are being shipped today with 65nm technology. Our nearest competitor is projecting its first 65nm product shipments at the end of 2006, whereas we started shipping 65nm products for revenue in October 2005.

People also need to realize that all 65nm processes are not equal. From what we've seen, Intel's transistors have the highest performance at a given leakage level and the lowest leakage at a given performance level. Also, Intel tends to have more advanced feature sizes; for example, the thinnest gate oxides and smallest SRAM cell sizes. As a historic note, the same objective benchmarks show that Intel transistors were superior to competitors' as we started volume production on 90nm and 130nm.

> "...We have the research discipline, the technology, and the process expertise to keep delivering industry-leading value for our customers now and in the future. That's always been our goal and our commitment."
>
> —Mark T. Bohr, Intel Senior Fellow and Director of Process Architecture and Integration

Benefits of 65nm Over 90nm Technology

Why should the average consumer care whether the processor they use is built on 65nm? What is the end-user benefit of this technology?

Intel's 65nm technology not only delivers the power and performance benefits that I talked about, but it effectively doubles the number of transistors on a chip, allowing Intel to pack hundreds of millions of transistors onto a single chip at ever-lower costs (per transistor). This capability gives us the foundation on which to deliver multi-core processors with more on-chip memory, and to design innovative features such as virtualization and security capabilities into future products.

Many of Intel's competitors talk about having 65nm in "pre-production," with plans to be in "production" shortly. Any thoughts on that?
In general, when our competitors say they are "in production," it means low-volume samples. When Intel is in production, it is high volume. As a testament to this, Intel's 65nm process has reached crossover with 90nm for CPU products, and this has occurred before our competitors have shipped any 65nm CPUs.

What are some of the special features of Intel's 65nm process?
This process implements second-generation strained silicon for improved performance and power characteristics. Our 65nm transistors have a more than 20 percent improvement in switching speed compared to 90nm at the same leakage level. Alternatively, transistor leakage can be reduced by around 5x at the same drive current (performance). The process also includes power-efficient interconnects, including an eighth layer of copper and second-generation carbon-doped oxide (CDO) dielectric.

Advanced and Future Intel Silicon Technologies

What is the difference between other manufacturers' strained silicon and Intel's?
A key difference is that the others have yet to put into production Intel's unique and highly beneficial use of SiGe source-drain structures to improve PMOS transistor performance, a feature that we've had in volume production since the end of 2003. Some others have tried to confuse the market with labels like "dual-strain" and "third-generation strained silicon," but their transistor technologies have yet to match the performance characteristics of our strained-silicon transistors.

What about other manufacturers' use of "silicon on insulator" (SOI)?
Intel's superior transistor performance and leakage have been achieved by meticulous engineering of "bulk CMOS" techniques, which offer the highest performance and best value to our customers. SOI adds substantial costs and complexity to the process. A comparison of published transistor switching speeds shows Intel's bulk transistors to have faster switching speeds at comparable leakage currents compared to the best published numbers on SOI transistors. A more troubling issue is that SOI has higher "thermal resistance" than bulk CMOS. As a result, SOI transistors are forced to run at temperatures higher than necessary. Higher operating temperatures could cause long-term reliability problems.

So does Intel rule out SOI completely?
No. There is a type of SOI called fully depleted (FD-SOI) that has merits beyond the partially depleted SOI that some of our competitors are using today. FD-SOI has been under active evaluation for some time at Intel.

What about the latest Intel announcements on tri-gate transistors with high-k/metal gate?
Our researchers have demonstrated CMOS tri-gate (3-D) transistors that are the first to integrate high-k gate dielectrics and strained silicon, producing record drive currents and transistor efficiency. (Read the Technology@Intel Magazine article, "Integrated CMOS Tri-Gate Transistors: Paving the way to Future Technology Generations," to learn more.) These transistors are an important option for achieving Intel's energy-efficient performance goals, and we believe that they could become the basic building block for future microprocessors sometime beyond the 45nm process technology node.

What other transistor technologies is Intel investigating?
There are many. As an example, at IEDM 2005, Intel disclosed that it had successfully integrated fully silicided (FUSI) gates and uniaxial strained silicon channels. The performance gains from strained channels and FUSI gates are fully additive. The result is the best transistor performance/leakage characteristics reported to date in the industry.

Intel Silicon Delivers Industry-Leading Value

So where does Intel stand today in the industry?

I think we're defining the leading edge in our labs and delivering it in our products. When looking at silicon process, it's important to separate the hype from the technology. Our well-known two-year process technology cycle is in cadence with our plan to introduce a new microarchitecture every two years.

Just look at 65nm. While everyone is talking about it, no company is producing 65nm chips in volume today, yet we have been doing so since October 2005. For another example, many of the fab automation features being hyped by our competitors were implemented at Intel more than 10 years ago, without fanfare. The bottom line is that we have the research discipline, the technology, and the process expertise to keep delivering industry-leading value for our customers now and in the future. That's always been our goal and our commitment.

More Info

You can discover much more about Intel's silicon technologies by visiting the following areas of the Intel Web site:

Energy-Efficient Performance
65-Nanometer Technology
Tri-Gate Transistors
45-Nanometer Technology

Mark T. Bohr Bio

Mark T. Bohr, Intel Senior Fellow, Technology and Manufacturing Group, Director of Process Architecture and Integration, Intel Corporation

Mark T. Bohr is an Intel Senior Fellow in the Technology and Manufacturing Group and Director of Process Architecture and Integration. He works in Intel's Logic Technology Development group located in Hillsboro, Oregon, where he is responsible for directing process development activities for Intel's advanced logic technologies.
He joined Intel in 1978 and has been responsible for process integration and device design on a variety of process technologies for dynamic RAM, static RAM, and microprocessor products. The technologies that he has helped to develop include: Intel's first CMOS technology in 1981, the world's first CMOS DRAM technology in 1983, Intel's first BiCMOS logic technology in 1992, and recent 90nm and 65nm logic technologies using strained silicon transistors and copper + low-k interconnects. He is currently directing development activities for Intel's 45nm logic technology.

Bohr received a B.S. in industrial engineering and an M.S.E.E., both from the University of Illinois, Urbana-Champaign. In 1998 he received the Distinguished Alumnus Award from the University of Illinois department of electrical and computer engineering. Bohr is a Fellow of the Institute of Electrical and Electronics Engineers and was the recipient of the 2003 IEEE Andrew S. Grove award. He presently serves on the executive committee for the VLSI Symposia. In 2005 he was elected to the National Academy of Engineering. He holds 31 patents in the area of integrated circuit processing and has authored or co-authored 39 published papers.

—End of Technology@Intel Magazine Article—